

Factorized Asymptotic Bayesian Inference for Factorial Hidden Markov Models

Shaohua Li

Nanyang Technological University
Singapore
shaohua@gmail.com

Ryohei Fujimaki

NEC Laboratories America
rfujimaki@nec-labs.com

Chunyan Miao

Nanyang Technological University
Singapore
ascymiao@ntu.edu.sg

Abstract

Factorial hidden Markov models (FHMMs) are powerful tools of modeling sequential data. Learning FHMMs yields a challenging simultaneous model selection issue, i.e., selecting the number of multiple Markov chains and the dimensionality of each chain. Our main contribution is to address this model selection issue by extending Factorized Asymptotic Bayesian (FAB) inference to FHMMs. First, we offer a better approximation of marginal log-likelihood than the previous FAB inference. Our key idea is to integrate out transition probabilities, yet still apply the Laplace approximation to emission probabilities. Second, we prove that if there are two very similar hidden states in an FHMM, i.e. one is redundant, then FAB will almost surely shrink and eliminate one of them, making the model parsimonious. Experimental results show that FAB for FHMMs significantly outperforms state-of-the-art nonparametric Bayesian iFHMM and Variational FHMM in model selection accuracy, with competitive held-out perplexity.

1 Introduction

The Factorial Hidden Markov Model (FHMM) [8] is an extension of the Hidden Markov Model (HMM), in which the hidden states are factorized into several independent Markov chains, and emissions (observations) are determined by their combination. FHMMs have found successful applications in speech recognition [16], source separation [18, 15, 11, 12], natural language processing [4] and bioinformatics [10]. A graphical model representation of the FHMM is shown in Fig.1.

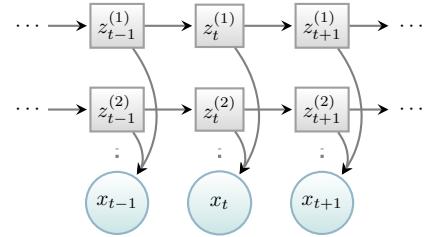


Figure 1: Graphical Model Representation of FHMMs.

Learning FHMMs naturally yields a challenging simultaneous model selection issue, i.e., how many independent Markov chains we need (layer-level model selection), and what is the dimensionality (number of hidden states) of each Markov chain. As the model space increases exponentially with the number of the Markov chains, it is not feasible to employ a grid search based method like cross-validation [2]. More sophisticated methods, like variational Bayesian inference (VB-FHMMs) [1] and the nonparametric Bayesian Infinite FHMMs (iFHMMs) [19] have been proposed, but they cannot fully address the issue. For example, iFHMMs restrict the dimensionality of each Markov chain to be binary. Further, high computational costs of VB-FHMMs and iFHMMs restrict their applicability to large scale problems.

This paper addresses the above model selection issue of FHMMs by extending Factorized Asymptotic Bayesian (FAB) inference, which is a recent model selection framework and has shown superior performance than nonparametric Bayesian inference for Mixture Models [7], Hidden Markov Models [6], Latent Feature Models [9], and Hierarchical Mixture of Experts Models [5]. Our method, namely FAB_{fhmm} , not only fully addresses the simultaneous model selection issue, but also offers the following two contributions.

1) Better marginal log-likelihood approximation with partial marginalization technique: Previous FAB inference for HMMs [6] has applied the Laplace approximations both to emission probabilities and to transition probabilities. Our approach applies

the Laplace approximations only to emission probabilities, and transition probabilities are integrated out. As we do not conduct asymptotic approximation but follow exact marginalization on transition probabilities, we can better approximate marginal log-likelihood.

2) A quantitative analysis of the shrinkage process: One of the strong features in FAB inference is a shrinkage effect, caused by FAB-unique asymptotic regularization, on hidden variables , i.e., redundant hidden states are automatically removed during the FAB EM-like iterative optimization. Although its strong model selection capability has been empirically confirmed [7, 6, 9], its mathematical behavior has not been well studied. This paper carefully investigates the shrinkage process, and proves that if there are two very similar hidden states in FHMMs, then one state would almost surely “die out”. We reveal the following chain reaction between E-steps and M-steps: if the variational probabilities of some states are shrunk in an E-step, then their corresponding model parameters will also be shrunk in the next M-step, causing these state to be shrunk further in future E-steps. Moreover, under certain condition, this shrinkage process is accelerating. This finding also partially answers the parameter identifiability problem of FHMMs.

2 Related Work

2.1 Factorial HMMs

The pioneering work of FHMMs was by Z. Ghahramani and M. Jordan [8]. The idea of FHMMs is to express emissions (observations) by combining M independent Markov chains. If the individual Markov chain have binary states, the FHMM can express 2^M different emission distributions. Their inference, *structured variational inference*, uses Baum-Welch algorithm [17] to learn parameters. Generally speaking, variational Bayesian (VB) inference [14, 1] can prune redundant latent states, but previous studies have suggested the pruning effect is not strong enough to achieve good model selection performance.

Recently, infinite FHMMs (iFHMMs) have been proposed [19] to address the model selection issue of FHMMs. iFHMMs employ the Markov Indian Buffet Process (mIBP) as the prior of the transition matrix of the infinite hidden states. Although iFHMMs offer strong model selection ability in learning FHMMs, they have a few limitations. First, iFHMMs restrict latent variables to be binary while the original FHMMs [8] have no such limitation. This restriction makes iFHMMs generate more represented states than variational methods, tending to overfit the data. In contrast, excluding such a restriction and allowing different Markov chains to have different numbers of hid-

den states may give us better understanding of the data. Second, the slice sampling used to optimize the iFHMM is considerably slower than variational methods, and therefore iFHMMs do not scale well to large scale scenarios.

2.2 FAB Inference

Factorized asymptotic Bayesian (FAB) inference has been originally proposed to address model selection of mixture models [7]. FAB inference selects the model which maximizes an asymptotic approximation of the marginal log-likelihood of the observed data, referred to as *Factorized Information Criterion* (FIC), with an EM-like iterative optimization procedure. Previous studies have extended FAB inference to HMMs (sequential) [6] and Latent Feature Models (factorial) [9], and have shown superiority against variational inference and nonparametric Bayesian methods in terms of model selection accuracy and computational cost. It is an interesting open challenge to investigate FAB inference for their intermediate (both sequential and factorial) models, i.e., FHMMs.

3 Factorial Hidden Markov Models

Suppose we have observed N independent sequences, denoted as $\mathbf{x}^N = \mathbf{x}^1, \dots, \mathbf{x}^N$. The n -th sequence is denoted as $\mathbf{x}^n = \mathbf{x}_1^n, \dots, \mathbf{x}_{T_n}^n$, where T_n is the length of the n -th sequence. Respectively, we denote the corresponding sequences of hidden state variables as $\mathbf{z}^N = \mathbf{z}^1, \dots, \mathbf{z}^N$, and each sequence $\mathbf{z}^n = \mathbf{z}_1^n, \mathbf{z}_2^n, \dots, \mathbf{z}_{T_n}^n$. \mathbf{z}_t^n consists of hidden variables of M independent HMMs: $\mathbf{z}_t^n = \mathbf{z}_t^{n,(1)}, \mathbf{z}_t^{n,(2)}, \dots, \mathbf{z}_t^{n,(M)}$. The m -th HMM has K_m hidden states, i.e., $\mathbf{z}_t^{n,(m)} \in \{0, 1\}^{K_m}$. We represent $\mathbf{z}_t^{n,(m)}$ as a $K_m \cdot 1$ binary vector, where the k -th component, $z_{t,k}^{n,(m)}$, denotes whether state $z_t^{n,(m)} = k$.

An FHMM model is specified as the following probability density:

$$p(\mathbf{x}^n, \mathbf{z}^n | \boldsymbol{\theta}) = \prod_{m=1}^M \left\{ p(\mathbf{z}_1^{n,(m)} | \boldsymbol{\alpha}^m) \prod_{t=2}^{T_n} p(\mathbf{z}_t^{n,(m)} | \mathbf{z}_{t-1}^{n,(m)}, \boldsymbol{\beta}^m) \right\} \prod_{t=1}^T p(\mathbf{x}_t^n | \mathbf{z}_t^n, \boldsymbol{\phi}), \quad (1)$$

where $\boldsymbol{\theta} = (\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\phi})$. The initial, transition and emission distributions are represented as $p(\mathbf{z}_1^{(m)} | \boldsymbol{\alpha}^m)$, $p(\mathbf{z}_t^{n,(m)} | \mathbf{z}_{t-1}^{n,(m)}, \boldsymbol{\beta}^m)$, and $p(\mathbf{x}_t^n | \mathbf{z}_t^n, \boldsymbol{\phi})$, respectively.

The m -th initial and transition distributions are defined as follows:

- $p(\mathbf{z}_1^{(m)} | \boldsymbol{\alpha}^{(m)}) = \prod_{k=1}^{K_m} (\alpha_k^{(m)})^{z_{1,k}^{(m)}};$
- $p(\mathbf{z}_t^{(m)} | \mathbf{z}_{t-1}^{(m)}, \boldsymbol{\beta}^{(m)}) = \prod_{j,k=1}^{K_m} (\beta_{j,k}^{(m)})^{z_{t-1,j} z_{t,k}^{(m)}}.$

Here $\boldsymbol{\alpha}^{(m)} = (\alpha_1^{(m)}, \dots, \alpha_{K_m}^{(m)})$, with $\sum_{k=1}^{K_m} \alpha_k^{(m)} = 1$; $\boldsymbol{\beta}^{(m)} = (\beta_{i,j}^{(m)})$, with each row $\boldsymbol{\beta}_k^{(m)}$ summing to 1.

By following the original FHMM work [8], this paper considers multidimensional Gaussian emission, while it is not difficult to extend our discussion to more general distributions like one in the exponential family. The emission distribution is jointly parameterized across the hidden states in the M HMMs as follows:

$$p(\mathbf{x}_t^n | \mathbf{z}_t^n, \phi) = \mathcal{N}(\mathbf{x}_t^n, \boldsymbol{\mu}_t^n, \mathbf{C}), \quad (2)$$

where $\boldsymbol{\mu}_t^n$ and \mathbf{C} are the mean vector and covariance matrix. A key idea (and difference from HMMs) in FHMMs is the construction of the mean vector $\boldsymbol{\mu}_t^n$. More specifically, the mean vector $\boldsymbol{\mu}_t^n$ is represented by the following linear combination:

$$\boldsymbol{\mu}_t^n = \sum_{m=1}^M \mathbf{W}^m \mathbf{z}_t^{n,(m)}, \quad (3)$$

where \mathbf{W}^m is a $D \times K_m$ matrix. The k -th row of \mathbf{W}^m is denoted as \mathbf{W}_k^m , specifying the contribution of the k -th state in the m -th HMM to the mean. Here, ϕ in (1) can be represented as $\phi = (\mathbf{W}, \mathbf{C})$.

4 Refined Factorized Information Criterion for FHMMs

FIC derivation starts from the following equivalent form of marginal log-likelihood:

$$\log p(\mathbf{x}^N | \mathcal{M}) \equiv \max_q \sum_{\mathbf{z}^N} q(\mathbf{z}^N) \log \left(\frac{p(\mathbf{x}^N, \mathbf{z}^N | \mathcal{M})}{q(\mathbf{z}^N)} \right) \quad (4)$$

\mathcal{M} represents a model and $\mathcal{M} = (M, K_1, \dots, K_M)$ in FHMMs. $q(\mathbf{z})$ is a variational distribution over the hidden states. For a state vector $\mathbf{z}_t^{n,m}$, its expectation under q , $\mathbf{E}_q[\mathbf{z}_t^{n,m}]$, is denoted in shorthand as $q(\mathbf{z}_t^{n,m})$, and the expectation of two consecutive states $\mathbf{E}_q[z_{t-1,j}^{n,m} \cdot z_{t,k}^{n,m}]$ as $q(z_{t-1,j}^{n,m}, z_{t,k}^{n,m})$.

A direct application of the technique proposed in [7, 6] leads to the following asymptotic approximation of the complete marginal log-likelihood:

$$\begin{aligned} p(\mathbf{x}^N, \mathbf{z}^N | \mathcal{M}) &= \prod_{m=1}^M \underbrace{\left\{ \int p(\mathbf{z}_1^{(m)} | \boldsymbol{\alpha}^m) p(\boldsymbol{\alpha}^m | \mathcal{M}) d\boldsymbol{\alpha}^m \right\}}_{\text{Laplace Approx.}} \\ &\quad \underbrace{\int \prod_{t=2}^{T_n} p(\mathbf{z}_t^{n,(m)} | \mathbf{z}_{t-1}^{n,(m)}, \boldsymbol{\beta}^m) p(\boldsymbol{\beta}^m | \mathcal{M}) d\boldsymbol{\beta}^m}_{\text{Laplace Approx.}} \\ &\quad \underbrace{\int \prod_{t=1}^T p(\mathbf{x}_t^n | \mathbf{z}_t^n, \phi) p(\phi | \mathcal{M}) d\phi}_{\text{Laplace Approx.}}. \end{aligned} \quad (5)$$

Instead, this paper proposes the FIC derivation with integrating out the initial and transition distributions.

$$\begin{aligned} p(\mathbf{x}^N, \mathbf{z}^N | \mathcal{M}) &= \prod_{m=1}^M \underbrace{\left\{ \int p(\mathbf{z}_1^{(m)} | \boldsymbol{\alpha}^m) p(\boldsymbol{\alpha}^m | \mathcal{M}) d\boldsymbol{\alpha}^m \right\}}_{\text{Integrated Out}} \\ &\quad \underbrace{\int \prod_{t=2}^{T_n} p(\mathbf{z}_t^{n,(m)} | \mathbf{z}_{t-1}^{n,(m)}, \boldsymbol{\beta}^m) p(\boldsymbol{\beta}^m | \mathcal{M}) d\boldsymbol{\beta}^m}_{\text{Integrated Out}} \\ &\quad \underbrace{\int \prod_{t=1}^T p(\mathbf{x}_t^n | \mathbf{z}_t^n, \phi) p(\phi | \mathcal{M}) d\phi}_{\text{Laplace Approx.}}. \end{aligned} \quad (6)$$

The original motivation of FIC is to approximate “intractable” Bayesian marginal log-likelihood. The key idea in (6) is to refine approximation by analytically solving the tractable integrations (w.r.t. $\boldsymbol{\alpha}^m$ and $\boldsymbol{\beta}^m$) and minimizing the asymptotic approximation error.

After integrating out all the parameters, and normalizing the Hessians w.r.t. \mathbf{W} and \mathbf{C}^{-1} , we obtain:

With uninformative conjugate priors on $\boldsymbol{\alpha}^m$ and $\boldsymbol{\beta}^m$, we can calculate (6) as follows:

$$\begin{aligned} &p(\mathbf{x}^N, \mathbf{z}^N | \mathcal{M}) \\ &\approx \prod_m \left(\frac{\prod_k \Gamma(c_{m,0,k} + 1)}{\Gamma(\sum_k c_{m,0,k} + K_m)} \prod_j \frac{\prod_k \Gamma(c_{m,j,k} + 1)}{\Gamma(\sum_k c_{m,j,k} + K_m)} \right) \\ &\quad \cdot \prod_{d=1}^D (2\pi)^{\frac{K_0+1}{2}} \left(\prod_{n=1}^N \frac{T_n}{2} \right)^{-\frac{1}{2}} c_d^{-\frac{K_0-2}{2}} |\hat{\mathcal{F}}_{\bar{\mathbf{W}}_d, \bar{c}_d}|^{-\frac{1}{2}} \\ &\quad \cdot \prod_{m,k=1}^{M,K_m} \left(\sum_{n,t=1}^{N,T_n} z_{t,k}^{n,(m)} \right)^{-\frac{1}{2}}, \end{aligned} \quad (7)$$

where $K_0 = \sum_m K_m$, and $\hat{\mathcal{F}}_{\bar{\mathbf{W}}_d, \bar{c}_d}$ is the normalized Hessian w.r.t. the d -th rows of \mathbf{W} and \mathbf{C}^{-1} at the maximum likelihood estimators $\bar{\mathbf{W}}, \bar{\mathbf{C}}$. The normalization makes $|\hat{\mathcal{F}}_{\bar{\mathbf{W}}_d, \bar{c}_d}| = O(1)$, by extracting the terms involving \mathbf{z} .

(4) computes the expectation of $\log p(\mathbf{x}^N, \mathbf{z}^N | \mathcal{M})$ w.r.t. the variational distribution q . Taking the logarithm of (7), many log-Gamma terms appear, whose exact expectations require to enumerate the exponentially many possible configurations of \mathbf{z} , which is infeasible. Thus we propose a first-order approximation to them, based on the following lemmas, where $\epsilon_1, \epsilon_2, \epsilon_3$ are small bounded errors.

Lemma 1. Suppose $\{z_1^n, \dots, z_{T_n}^n\}_{n=1}^N$ are N sequences of Bernoulli random variables. Within each sequence, $\{z_t^n\}_{t=1}^{T_n}$ are independent with each other. Let $y_n = \sum_i z_i^n$, $\bar{y}_n = E[y_n]$. Besides, there are N

numbers $\{\hat{y}_n\}$, $\forall k, \hat{y}_n \approx \bar{y}_n$. When all T_n are large enough:

- 1) $\mathbf{E}[\log \Gamma(y_n)] = y_n \log \bar{y}_n - (\bar{y}_n + \frac{1}{2} \log \bar{y}_n) + \log 2\pi + \epsilon_1$.
- 2) $\mathbf{E}[\sum_n \log \Gamma(y_n) - \log \Gamma(\sum_n y_n)] = \sum_n y_n \log \left(\frac{\bar{y}_n}{\sum_{m=1}^N \hat{y}_m} \right) + \frac{1}{2} \log \left(\sum_n \hat{y}_n \right) - \frac{1}{2} \sum_n \log \hat{y}_n + (N-1) \log 2\pi + \epsilon_2$.
- 3) $\mathbf{E}[\log y_n] = \log \hat{y}_n + \frac{1}{\hat{y}} (\bar{y}_n - \hat{y}_n) + \epsilon_3$.

Proof can be found in the Appendix.

As defined above, $c_{m,j,k} = \sum_{n,t=1}^{N,T_n-1} z_{t,j}^{n,m} z_{t+1,k}^{n,m}$. Since $\{z_{t,j}^{n,m}\}_{t=1}^{T_n}$ follows a Markov process, the dependency between $z_{t,j}^{n,m}$ and $z_{t+\Delta t,j}^{n,m}$ vanishes quickly when the time gap Δt increases, leading to the uncoupling between $z_{t,j}^{n,m} z_{t+1,k}^{n,m}$ and $z_{t+\Delta t,j}^{n,m} z_{t+1+\Delta t,k}^{n,m}$. Thus we can approximately regard $c_{m,j,k}$ as the sum of independent Bernoulli random variables, and therefore Lemma 1 apply. The same argument applies to $c_{m,0,k}$.

In order to avoid recursion relations w.r.t. q during the inference, we introduce an auxiliary distribution \hat{q} , which is always close to q , in that $\hat{q}(z_{t,k}^{n,m}) \approx q(z_{t,k}^{n,m}), \hat{q}(z_{t,j}^{n,m}, z_{t+1,k}^{n,m}) \approx q(z_{t,j}^{n,m}, z_{t+1,k}^{n,m})$. The Gamma terms in (7) will be expanded about the expectations of their parameters w.r.t. \hat{q} . We will discuss how to choose \hat{q} in Section 5.2.

Taking the logarithm of (7), applying Lemma 1, and dropping asymptotically small terms (including the determinants of the normalized Hessian matrices), we derive the approximation of $p(\mathbf{x}^N, \mathbf{z}^N | \mathcal{M})$. Plugging it into (4), we obtain an asymptotic approximation of the lower bound of $\log p(\mathbf{x}^N | M)$:

$$\begin{aligned} & \log p(\mathbf{x}^N | \mathcal{M}) \geq \mathcal{J}(q, \mathbf{x}^N) \\ & \approx \sum q(\mathbf{z}^N) \left[\sum_{n,t} \left(\log p(\mathbf{x}_t^n | \mathbf{z}_t^n, \bar{\mathbf{W}}, \bar{\mathbf{C}}) + \sum_{m,k} z_{t,k}^{n,m} \log \delta_k^m \right) \right. \\ & \quad \left. + \sum_{m,n=1}^{M,N} \left(\sum_k z_{0,k}^{n,m} \log \hat{\alpha}_k^m + \sum_{t=1}^{T_n-1} \sum_{j,k} z_{t,j}^{n,m} z_{t+1,k}^{n,m} \log \hat{\beta}_{j,k}^m \right) \right. \\ & \quad \left. - \log q(\mathbf{z}^N) \right] + \sum_{m,n,t} \log \Delta^m - \frac{D}{2} \sum_{m,k=1}^{M,K_m} \log(\hat{c}_{m,\cdot,k}) \\ & \quad - \frac{1}{2} \sum_{m,k} \log(\hat{c}_{m,0,k} + 1) - \frac{1}{2} \sum_{m,j,k} \log(\hat{c}_{m,j,k} + 1) \\ & \quad + \frac{1}{2} \sum_{m,j} \log \left(\sum_k \hat{c}_{m,j,k} + K_m \right) + \frac{1}{2} \sum_m \log(N + K_m) \\ & \quad + \frac{D}{2} \sum_m K_m + \sum_m (K_m^2 - 1) \log 2\pi + \epsilon, \end{aligned} \quad (8)$$

with the definitions

$$\begin{aligned} \hat{q}(z_{t,i}^{n,m}) &= E_q(z_{t,i}^{n,m}), \quad q(z_{t,j}^{n,m}, z_{t+1,k}^{n,m}) = E_q(z_{t,j}^{n,m}, z_{t+1,k}^{n,m}), \\ \hat{c}_{m,0,k} &= \sum_{n=1}^N \hat{q}(z_{1,k}^{n,m}), \quad \hat{c}_{m,\cdot,k} = \sum_{n,t=1}^{N,T_n} \hat{q}(z_{t,k}^{n,m}), \\ \hat{c}_{m,j,k} &= \sum_{n,t=1}^{N,T_n-1} \hat{q}(z_{t,j}^{n,m}, z_{t+1,k}^{n,m}), \\ \delta_k^m &= \frac{1}{\Delta^m} \exp \left\{ - \frac{D}{2 \sum_{n,t=1}^{N,T_n} \hat{q}(z_{t,k}^{n,m})} \right\}, \\ \hat{\alpha}_k^m &= \frac{1 + \sum_{n=1}^N \hat{q}(z_{1,k}^{n,m})}{K_m + N}, \\ \hat{\beta}_{j,k}^m &= \frac{1 + \sum_{n,t=1}^{N,T_n-1} \hat{q}(z_{t,j}^{n,m}, z_{t+1,k}^{n,m})}{K_m + \sum_{n,t=1}^{N,T_n-1} \hat{q}(z_{t,j}^{n,m})}, \end{aligned} \quad (9)$$

and Δ^m is a normalization constant that makes $\sum_{k=1}^{K_m} \delta_k^m = 1$, ε is a small constant error bound of the approximations. Note the ML estimators $\bar{\mathbf{W}}, \bar{\mathbf{C}}$ inside the summation are subject to the specific instantiation of \mathbf{z}^N .

In (9), $\hat{\alpha}$ and $\hat{\beta}$ can be viewed as the estimated initial transition probabilities of the FHMM w.r.t. \hat{q} .

The FIC of the FHMM, denoted by FIC_{fhmm} , is obtained by maximizing $\mathcal{J}(q, \mathbf{x}^N)$ w.r.t. q :

$$\text{FIC}(\mathbf{x}^N, \mathcal{M}) = \max_q \{\mathcal{J}(q, \mathbf{x}^N)\}. \quad (10)$$

Similar to FIC_{mm} and FIC_{hmm} , the use of FIC_{fhmm} as the approximation of the observed data log-likelihood under a certain model is justified:

Theorem 2. $\text{FIC}(\mathbf{x}^N, \mathcal{M})$ is asymptotically equivalent to $\log p(\mathbf{x}^N | \mathcal{M})$.

The proof is analogous to that of FIC_{mm} and FIC_{hmm} [7, 6], and omitted here.

5 FAB for FHMMs

5.1 FAB's Lower Bound of FIC

Since $\bar{\mathbf{W}}, \bar{\mathbf{C}}$ depends on the specific instantiation of \mathbf{z}^N , in order to compute FIC_{fhmm} exactly, we need to evaluate $\bar{\mathbf{W}}, \bar{\mathbf{C}}$ for each instantiation of \mathbf{z}^N , which is infeasible. So we bound $\text{FIC}(\mathbf{x}^N, M)$ from below by setting all $\bar{\mathbf{W}}, \bar{\mathbf{C}}$ to the same values, and get a relaxed lower bound \mathcal{G} :

$$\mathcal{G}(q, \mathbf{x}^N, \mathbf{W}, \mathbf{C}) < \mathcal{J}(q, \mathbf{x}^N) \leq \log p(\mathbf{x}^N | \mathcal{M}).$$

FAB_{fhmm} is an optimization procedure that tries to maximize the above lower bound of FIC_{fhmm} :

$$\mathcal{M}^*, q^*, \mathbf{W}^*, \mathbf{C}^* = \underset{\mathcal{M}, q, \mathbf{W}, \mathbf{C}}{\operatorname{argmax}} \mathcal{G}(q, \mathbf{x}^N, \mathbf{W}, \mathbf{C}). \quad (11)$$

5.2 FAB Variational EM Algorithm

Let us first fix the structural parameters of the FHMM model, i.e. M, K_1, \dots, K_M . Then our objective is to optimize (11) w.r.t. $(q, \mathbf{W}, \mathbf{C})$. In this phase FAB is a typical variational EM Algorithm, iterating between E-steps and M-steps. We denote the i -th iteration with the superscript $\{i\}$.

In (8), the auxiliary distribution \hat{q} is required to be close to q . So during the FAB optimization, we always set \hat{q} to be the variational distribution in the previous iteration, i.e. $\hat{q}^{\{i\}} = q^{\{i-1\}}$.

In the i -th E-step, we fix $\mathbf{W}, \mathbf{C} = \mathbf{W}^{\{i-1\}}, \mathbf{C}^{\{i-1\}}$, and $\hat{q} = q^{\{i-1\}}$. The E-step computes the q which maximizes (11).

The exact E-step for FHMMs is intractable [8], and thus in E-step, we adopt the *Mean-field Variational Inference* proposed in [8].

5.2.1 FAB E-Step

Mean-Field Variational Inference The *Mean-Field Variational Inference* approximates the FHMM with M uncoupled HMMs, by introducing a variational parameter $\mathbf{h}_t^{n,(m)}$ for each layer variable $\mathbf{z}_t^{n,(m)}$. $\mathbf{h}_t^{n,(m)}$ approximates the contribution of $\mathbf{z}_t^{n,(m)}$ to the corresponding observation \mathbf{x}_t^n .

The structured variational distribution q is defined as

$$q(\mathbf{z}^N | \mathbf{h}) = \frac{1}{Z} \prod_{n,m=1}^{N,M} \left\{ \hat{q}(\mathbf{z}_1^{n,(m)} | \mathbf{h}) \prod_{t=2}^{T_n} \hat{q}(\mathbf{z}_t^{n,(m)} | \mathbf{z}_{t-1}^{n,(m)}, \mathbf{h}) \right\}, \quad (12)$$

where Z is the normalization constant, and

$$\begin{aligned} \hat{q}(\mathbf{z}_1^{n,(m)} | \mathbf{h}) &= \prod_{k=1}^{K_m} (h_{1,k}^{n,(m)} \hat{\alpha}_k^{n,(m)})^{z_{1,k}^{n,(m)}}, \\ \hat{q}(\mathbf{z}_t^{n,(m)} | \mathbf{z}_{t-1}^{n,(m)}, \mathbf{h}) &= \prod_{k=1}^{K_m} \left(h_{t,k}^{n,(m)} \prod_{j=1}^{K_m} (\hat{\beta}_{j,k}^{(m)})^{z_{t-1,j}^{n,(m)}} \right)^{z_{t,k}^{n,(m)}}. \end{aligned} \quad (13)$$

Note $\mathbf{h}_t^{n,(m)}$ is a $1 \times K_m$ vector, which gives a bias for each of the K_m settings of $z_{t,k}^{n,(m)}$.

We obtain a system of equations for $\mathbf{h}_t^{n,(m)}$ and $q(\mathbf{z}_t^{n,(m)})$, which minimize $\text{KL}(q || \hat{p}(\mathbf{z}^N | \mathbf{x}^N, \boldsymbol{\theta}))$, by setting its derivation w.r.t. $\mathbf{h}_t^{n,(m)}$ to 0:

$$\mathbf{h}_t^{n,(m)} = \text{diag}\{\boldsymbol{\delta}^m\} \exp \left\{ \mathbf{W}^{m'} \mathbf{C}^{-1} \tilde{\mathbf{x}}_t^{n,(m)} - \frac{1}{2} \Lambda^m \right\}, \quad (14)$$

where Λ^m is the vector consisting of the diagonal elements of $\mathbf{W}^{m'} \mathbf{C}^{-1} \mathbf{W}^m$; $\text{diag}(\mathbf{v})$ is an operator that

constructs a matrix whose diagonal is \mathbf{v} , and all off-diagonal elements are 0; $\tilde{\mathbf{x}}_t^{n,(m)}$ is the residual:

$$\tilde{\mathbf{x}}_t^{n,(m)} = \mathbf{x}_t^n - \sum_{l \neq m}^M W^l q(\mathbf{z}_t^{n,(l)}). \quad (15)$$

In (14), \mathbf{h} depends on q , but q also depends on \mathbf{h} in the Forward-Backward routine, in a complicated way. Such an interdependence makes the exact solution difficult to find. Therefore in (15), we use $q^{\{i-1\}}$ to get approximate solutions of \mathbf{h} .

Note the bias $h_{t,k}^{n,(m)}$ incorporates the effects of the shrinkage factor δ_k^m . Through the bias, the shrinkage factors will induce the *Hidden State Shrinkage* to be discussed later.

Forward-Backward Routine After updating the variational parameters \mathbf{h} , we use the Forward-Backward algorithm to compute the sufficient statistics $q(\mathbf{z}_t^{n,(i)})$, and $q(\mathbf{z}_{t-1}^{n,(i_1)} \mathbf{z}_t^{n,(i_2)'}).$

The following recurrence relations for the forward quantities $f_{t,k}^{n,(m)}$ and backward quantities $b_{t,k}^{n,(m)}$ is derived from the definition (12-13) of q :

$$\begin{aligned} f_{t,k}^{n,(m)} &= \begin{cases} \frac{1}{\zeta_1^{n,(m)}} h_{1,k}^{n,(m)} \hat{\alpha}_k^{(m)} & \text{if } t=1 \\ \frac{1}{\zeta_t^{n,(m)}} h_{t,k}^{n,(m)} \sum_{j=1}^{K_m} f_{t-1,j}^{n,(m)} \hat{\beta}_{j,k}^{(m)} & \text{if } t>1 \end{cases} \\ b_{t,k}^{n,(m)} &= \begin{cases} \frac{1}{\zeta_{t+1}^{n,(m)}} \sum_{j=1}^{K_m} \hat{\beta}_{k,j}^{(m)} h_{t+1,j}^{n,(m)} b_{t+1,j}^{n,(m)} & \text{if } t < T_n \\ 1 & \text{if } t = T_n \end{cases} \end{aligned} \quad (16)$$

where $\zeta_t^{n,(m)}$ is the normalization constant that makes $\sum_{k=1}^{K_m} f_{t,k}^{n,(m)} = 1$.

The sufficient statistics are computed thereafter:

$$q(\mathbf{z}_{t,k}^{n,(m)}) = f_{t,k}^{n,(m)} b_{t,k}^{n,(m)}, \quad (17)$$

$$q(\mathbf{z}_{t-1,j}^{n,(m)}, \mathbf{z}_{t,k}^{n,(m)}) = \frac{1}{\zeta_t^{n,(m)}} f_{t-1,j}^{n,(m)} \hat{\beta}_{j,k}^{(m)} h_{t,k}^{n,(m)} b_{t,k}^{(m)}. \quad (18)$$

5.2.2 M-Step

In the M-step, we fix $q = \hat{q} = q^{\{i\}}$, and maximize (11) w.r.t. \mathbf{W}, \mathbf{C} , obtaining their optimal values:

$$\begin{aligned} \mathbf{C} &= \text{mdig} \left\{ \frac{1}{\sum_{n,t=1}^{N,T_n}} \sum_{n,t=1}^{N,T_n} \left(\mathbf{x}_t^n \mathbf{x}_t^{n'} - \sum_{m=1}^M W^m q(\mathbf{z}_t^{n,(m)}) \mathbf{x}_t^{n'} \right) \right\} \\ \mathbf{W} &= \left(\sum_{n,t=1}^{N,T_n} \mathbf{x}_t^n q(\mathbf{z}_t^{n'}) \right) \left(\sum_{n,t=1}^{N,T_n} \mathbf{E}_q [\mathbf{z}_t^n \mathbf{z}_t^{n'}] \right)^\dagger, \end{aligned} \quad (19)$$

where: 1) In the equation of \mathbf{W} , \dagger is the Moore-Penrose pseudo-inverse, $\mathbf{W} = (\mathbf{W}^1 \cdots \mathbf{W}^M)$, \mathbf{z}_t^n is an $(\sum_{m=1}^M K_m) \times 1$ vector of the concatenation of $\mathbf{z}_t^{n,(1)}, \dots, \mathbf{z}_t^{n,(M)}$; 2) In the equation of \mathbf{C} , the $\text{mdia}(X)$ operator sets all the off-diagonal elements of matrix X to 0, leaving the diagonal ones intact.

5.3 Automatic Hidden State Shrinkage

The shrinkage factors δ_k^m induce the *hidden state shrinkage* effect, which is a central mechanism to FAB_{fhhmm}'s model selection ability. Previous FAB methods [7, 6, 9] report similar effects.

In the beginning, we initialize the model with a large enough number M of layers, and all layers with a large enough number K_{\max} of hidden state values. During the EM iterations, suppose there are two similar hidden states i, j in the m -th layer, and state i is less favored than state j , i.e. $\sum_t q(z_{t,i}^{(m)}) < \sum_t q(z_{t,j}^{(m)})$. According to (9), $\delta_i^m < \delta_j^m$. In the next E-step, the variational parameters of the i -th state $\{h_{t,i}^{(m)}\}_t$ are more down-weighted than those of the j -th state $\{h_{t,j}^{(m)}\}_t$, according to (14). From the Forward-Backward update equations (16) one can see $q(z_i^{(m)})$ decreases with $h_{t,i}^{(m)}$. Therefore $\{q(z_i^{(m)})\}$ decreases more than $\{q(z_j^{(m)})\}$, which in turn causes smaller transition probabilities into state i in the next M-step. Subsequently, smaller transition probabilities into state i make $\sum_t q(z_{t,i}^{(m)})$ and δ_i^m even smaller compared to $\sum_t q(z_{t,j}^{(m)})$ and δ_j^m , respectively. This process repeats and reinforces itself, making state i less and less important. When $\sum_t q(z_{t,i}^{(m)}) < L$, a pre-specified threshold, then FAB_{fhhmm} regards state i as redundant and removes it. If all but one states of a layer are removed, then we know this layer is redundant and FAB_{fhhmm} removes this layer. When FAB_{fhhmm} converges, we obtain a more parsimonious model with the number of layers M' , the number of hidden states in each layers $K_1^*, \dots, K_{M'}^*$, the optimal model parameters $\mathbf{W}^*, \mathbf{C}^*$, and the variational distribution q^* .

This intuition is formulated in the following theorem.

Theorem 3. Suppose we have one sufficiently long training sequence of observations $\mathbf{x}_T = \mathbf{x}_1, \dots, \mathbf{x}_T, T \gg 1$. In the m -th layer, two states i, j are initialized with proportional initial and inwards transition probabilities, i.e., $\hat{\alpha}_i = \rho_0 \hat{\alpha}_j$, and $\forall k, \hat{\beta}_{k,i} = \rho_0 \hat{\beta}_{k,j}, \rho < 1$, and identical outwards transition probabilities $\forall k, \hat{\beta}_{i,k} = \hat{\beta}_{j,k}$. Then: 1) after the first EM iteration, the two states have identical weight vectors and outwards probabilities: $W_i^m = W_j^m, \forall k, \hat{\beta}'_{i,k} = \hat{\beta}'_{j,k}$, and proportional inwards proba-

bilities: $\hat{\alpha}'_i = \rho_1 \hat{\alpha}'_j, \forall k, \hat{\beta}'_{k,i} = \rho_1 \hat{\beta}'_{k,j}$, but $\rho_1 < \rho_0$. This trend preserves in all the following iterations. 2) When the probabilities of state j stabilize, this shrinkage process on state i is accelerating: after the n -th iteration, the probability ratio ρ_n satisfy: $\rho_n / \rho_{n-1} < \rho_{n-1} / \rho_{n-2}$.

Proof. For notation simlicity, we drop the superscript n for the sequence number.

1) For initialization all variational parameters \mathbf{h} are set to 1. From the Forward-Backward recurrence relations (16), it is easily seen that $f_{t,i}^{(m)} = \rho_0 f_{t,j}^{(m)}, b_{t,i}^{(m)} = b_{t,j}^{(m)}$. Thus from (17),(18), the sufficient statistics of the variational probability q satisfy, for $\forall t$:

$$q(z_{t,i}^{(m)}) = \rho_0 q(z_{t,j}^{(m)}), q(z_{t-1,i}^{(m)} z_{t,k}^{(m)}) = \rho_0 q(z_{t-1,j}^{(m)} z_{t,k}^{(m)}) \quad (20)$$

In the M-step, when updating \mathbf{W} , one relevant equation is

$$\sum_t \mathbf{x}_t q(\mathbf{z}_t^m)' = \sum_{l \neq m} \mathbf{W}^l \sum_t q(\mathbf{z}_t^l) q(\mathbf{z}_t^m)' + \mathbf{W}^m \sum_t \text{diag}\{q(\mathbf{z}_t^m)\} \quad (21)$$

We single out the i -th and j -th columns of both sides of (21):

$$\begin{aligned} \sum_t q(z_{t,i}^{(m)}) \mathbf{x}_t &= \sum_{l \neq m} \mathbf{W}^l \sum_t q(z_{t,i}^{(m)}) q(\mathbf{z}_t^l)' + \sum_t \text{diag}\{q(z_{t,i}^{(m)})\} \mathbf{W}_i^m \\ \sum_t q(z_{t,j}^{(m)}) \mathbf{x}_t &= \sum_{l \neq m} \mathbf{W}^l \sum_t q(z_{t,j}^{(m)}) q(\mathbf{z}_t^l)' + \sum_t \text{diag}\{q(z_{t,j}^{(m)})\} \mathbf{W}_j^m \end{aligned} \quad (22)$$

By plugging (20) into (22), we find $\mathbf{W}_i^m = \mathbf{W}_j^m$.

In the next E-step, using (14), we can identify

$$\begin{aligned} \frac{h_{t,i}^{n,(m)}}{h_{t,j}^{n,(m)}} &= \frac{\delta_i^m}{\delta_j^m} = \exp\left\{\frac{D}{2 \sum_t q(z_{t,j}^{n,m})} - \frac{D}{2 \sum_t q(z_{t,i}^{n,m})}\right\} \\ &= \exp\left\{-\left(\frac{1-\rho_0}{\rho_0}\right) \frac{D}{2 \sum_t q(z_{t,j}^{n,m})}\right\}, \end{aligned} \quad (23)$$

which is constant for any t . We denote it as λ_0 . Since $\rho_0 < 1$, obviously $\lambda_0 < 1$.

In the Forward-Backward routine, by examining equations (16), one can see $f_{t,i}^{(m)} = \rho_0 \lambda_0 f_{t,j}^{(m)}, b_{t,i}^{(m)} = b_{t,j}^{(m)}$. Let $\rho_1 = \rho_0 \lambda_0$. Again from (17),(18), the sufficient statistics of the new variational distribution q' satisfy, for $\forall t$:

$$\begin{aligned} q'(z_{t,i}^{(m)}) &= \rho_1 q'(z_{t,j}^{(m)}), \\ q'(z_{t-1,i}^{(m)} z_{t,k}^{(m)}) &= \rho_1 q'(z_{t-1,j}^{(m)} z_{t,k}^{(m)}), \\ q'(z_{t,k}^{(m)} z_{t+1,i}^{(m)}) &= \rho_1 q'(z_{t,k}^{(m)} z_{t+1,j}^{(m)}). \end{aligned}$$

When we update $\hat{\alpha}$ and $\hat{\beta}$, since the training sequence is sufficiently long ($T \gg 1$), the pseudocount “1” and “ K_m ” in their update equations can be ignored. Therefore one can easily obtain

$$\begin{aligned}\hat{\alpha}'_i &= \rho_1 \hat{\alpha}'_j, \\ \forall k, \hat{\beta}'_{k,i} &= \rho_1 \hat{\beta}'_{k,j}, \\ \forall k, \hat{\beta}'_{i,k} &= \hat{\beta}'_{j,k}.\end{aligned}$$

Obviously $\rho_1 < \rho_0$. That is, the new ratio of the initial/inwards probabilities becomes smaller, while the outwards probabilities remain unchanged.

The same argument holds for all subsequent iterations.

2) In the n -th iteration, let the ratio in (23) be denoted as λ_n . From the above arguments, $\rho_n/\rho_{n-1} = \lambda_n$, thus it is sufficient to prove $\lambda_n < \lambda_{n-1}$.

If state j “survives” the shrinkage, then after many iterations its probabilities would become stabilized, i.e., $\sum_t q(z_{t,j}^{n,m})$ is almost the same between the n -th and $(n-1)$ -th iterations. In this condition, λ_n decreases with ρ_n , together with $\rho_n < \rho_{n-1}$, one can conclude $\lambda_n < \lambda_{n-1}$. \square

Theorem 5 sheds light on FAB_{*fhmm*}’s shrinkage mechanism: given two similar components, if they are not initialized evenly (which is almost always the case due to randomness), then they will “compete” for probability mass, until one dominates the other. Another important observation is, the M-step “relays” the shrinkage effect between consecutive E-steps, and thus running multiple E-steps in a row would not gain much speed-up as observed on Latent Feature Models [9].

5.4 Parameter Identifiability

Although HMMs/FHMMs are identifiable in the level of equivalence classes [13], generally they are not identifiable if two states have very similar outwards transition probabilities and emission distributions. Formally, if states i, j satisfy: $\forall k, \beta_{i,k} \approx \beta_{j,k}$, and $p(x|z=i) \approx p(x|z=j)$, then we can alter the inwards transition probabilities of i, j without changing the probability law of this model, as long as the sum of their inwards transition probabilities is fixed, i.e., $\hat{\beta}_{\cdot,i} + \hat{\beta}_{\cdot,j} = \beta_{\cdot,i} + \beta_{\cdot,j}$. This forms an equivalence class with infinitely many parameters settings (one extreme is state i is absorbed into j), all of which fit the observed data equally well. Thus traditional point estimation methods are unable to compare these parameter settings and pick better ones.

In contrast, Theorem 5 reveals that, FAB_{*fhmm*} will almost surely shrink and remove one of such two states, picking a parsimonious parameter setting (e.g. the one with i absorbed into j) among this equivalence class.

6 Synthetic Experiments

To test our inference algorithm, we ran experiments on a synthetic data set. The performance of FAB_{*fhmm*} (denoted as “RFAB”, i.e. *Refined* FAB) and three major competitors, i.e., conventional FAB without marginalization (denoted as “FAB”), variational Bayesian FHMM (denoted as “VB”), and iFHMM, was compared.

6.1 Experimental Settings

We constructed a ground truth model FHMM, which has 3 layers, with state numbers $(2, 2, 3)$. These states comprise a total state space of $2 \cdot 2 \cdot 3 = 12$ combinations. The observations are 3-dimensional, and the covariance matrix $\Sigma = \begin{pmatrix} 0.4 & 0 & 0 \\ 0 & 0.4 & 0 \\ 0 & 0 & 0.4 \end{pmatrix}$. The three mean matrices $\{\mathbf{W}^m\}$, initial probabilities $\{\boldsymbol{\alpha}^m\}$ and transition matrices $\{\boldsymbol{\beta}^m\}$ were randomly generated, and omitted here.

Two sequences of length $T = 2000$ were randomly generated. One sequence was used for training, and the other was used for testing.

RFAB, FAB and VB-FHMM are all initialized with 3 HMMs, and 10 states in each HMM. Compared to the true model, this setting has a lot of redundant states. For VB-FHMM, we used the “component death” idea in [1], i.e. if a hidden state receives too little probability mass, then it will be removed. This scheme is similar to FAB’s pruning scheme, despite the fact that the inference algorithm does not incorporate the shrinkage regularization.

We obtained the Matlab code of iFHMM from Van Gael. We used its default hyperparameters, i.e.: $\alpha \sim \Gamma(1 + K, 1 + H_T)$, where K is the number of hidden states, and H_T is the T -th harmonic number; $\sigma_Y, \sigma_W \sim \Gamma(1, 1)$. In order to see how iFHMM performs the model selection on different hidden state spaces, we initialized iFHMM with hidden state numbers $K \in \{4, 7, 10, 15\}$.

The maximal iterations of all models other than iFHMM are set to 1000. iFHMM is set to have 500 iterations for burn-in and 500 iterations for sampling. Each algorithm was tested for 10 trials.

6.2 Results and Discussions

Three indices of the models were collected: 1) the eventually obtained state numbers; 2) the log-likelihood of the training sequence; and 3) the predictive log-likelihood of the test sequence.

The state number in the three HMMs are first sorted from smallest to largest, and then averaged. The av-

	K_1	K_2	K_3
RFAB	1.9 (0.4)	3.0 (0.6)	4.8 (1.1)
FAB	1.4 (0.6)	2.2 (0.7)	4.0 (1.5)
VB	3.2 (1.5)	4.7 (1.2)	6.3 (1.8)
iFHMM/4	4.4 (0.3)	/	/
iFHMM/7	7.3 (0.7)	/	/
iFHMM/10	9.5 (1.2)	/	/
iFHMM/15	13.6 (1.3)	/	/

Table 1: Average State Numbers

	Training	Testing
RFAB	-10937 (254)	-12416 (425)
FAB	-11848 (322)	-13647 (598)
VB	-11239 (336)	-14430 (662)
iFHMM/4	-12256 (293)	-14376 (329)
iFHMM/7	-13471 (274)	-14687 (385)
iFHMM/10	-13949 (214)	-15116 (302)
iFHMM/15	-15514 (201)	-16895 (342)

Table 2: Training and Testing Data Log-likelihoods

erage values, and their standard deviations (in parentheses) are reported in the following table:

From Table 1, we can see that FAB obtained the most parsimonious learned models, but often it removes too many states, such that one HMM has only one state left. In contrast, RFAB is more “conservative”, as it keeps more states. This difference is probably because after marginalizing α and β , the pseudocount “1” appears in the update equations 9 of $\hat{\alpha}$ and $\hat{\beta}$, which smooths their estimated values. VB-FHMM removed some states, but there are still a few redundant states left. It indicates that without the shrinkage regularization, the model selection ability is limited. iFHMM almost always stays around its initial number of hidden states, no matter how many states they are initialized with. This indicates the model selection ability of iFHMM is weak.

The training and testing data log-likelihoods are compared in Table 2. RFAB achieves the best log-likelihoods on both sequences, followed by FAB. VB-FHMM achieves better log-likelihood than FAB on the training sequence, but worse on the test sequence, which suggests that it overfits the training data. iFHMM’s performance is significantly inferior in general, and degrades quickly as the initial K increases.

Fig.2 shows how the hidden state numbers K change over iterations in a typical trial. The final state numbers at convergence are shown at the end of each line. RFAB and FAB converge much faster than other

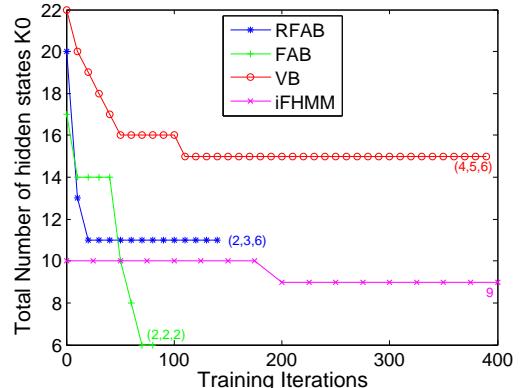


Figure 2: State Number Evolves over Iterations

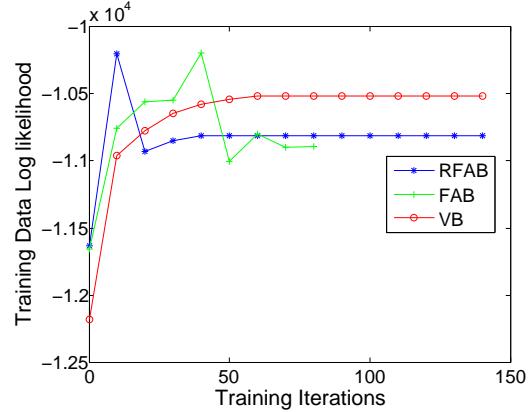


Figure 3: Training Data Log-likelihood

methods (< 200 iterations), which shows the shrinkage regularization accelerates the pruning of redundant states.

Fig.3 shows how training data log-likelihood changes over iterations in the same trial as in Fig.2. The log-likelihoods of RFAB and FAB on the training data are quite close. The log-likelihoods of iFHMM are not shown in this figure, as they are well below other methods (around -13500).

7 Conclusions and Future Work

The FHMM is a flexible and expressive model, but it also poses a difficult challenge on how to set its free, structural parameters. In this paper, we have successfully extended the recently-developed FAB framework onto the FHMM to address its model selection problem. We have derived a better asymptotic approximation of the data marginal likelihood than conventional FAB, by integrating out the initial and transition probabilities. Based on this refined marginal likelihood, an EM-like iterative optimization procedure,

namely FAB_{fhmm} , has been developed, which can find both good model structures and model parameters at the same time. Experiments on a synthetic data set have shown that FAB_{fhmm} obtains more parsimonious and better-fit models than the state-of-the-art non-parametric iFHMM and variational FHMM.

In addition, we have proved that, during the FAB_{fhmm} 's shrinkage process on hidden variables, if there are two very similar hidden states, then one state would almost surely “die out” and be pruned. This proof theoretically consolidates the shrinkage process which we observe in experiments.

References

- [1] M. Beal. *Variational algorithms for approximate Bayesian inference*. PhD thesis, University of London, 2003. 1, 2.1, 6.1
- [2] G. Celeux and J.-B. Durand. Selecting hidden markov model state number with cross-validated likelihood. *Computational Statistics*, 23:541–564, 2008. 1
- [3] M. T. Chao and W. E. Strawderman. Negative moments of positive random variables. *Journal of the American Statistical Association*, 67(338):429–431, 1972. A
- [4] K. Duh. Jointly labeling multiple sequences: a factorial hmm approach. In *Proceedings of the ACL Student Research Workshop*, ACLstudent '05, pages 19–24, Stroudsburg, PA, USA, 2005. Association for Computational Linguistics. 1
- [5] R. Eto, R. Fujimaki, S. Morinaga, and T. Hiroshi. Fully-automatic bayesian piece-wise sparse linear models. In *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2014. 1
- [6] R. Fujimaki and K. Hayashi. Factorized asymptotic bayesian hidden markov models. In *ICML*, volume 22, pages 400–408, 2012. 1, 2.2, 4, 4, 5.3
- [7] R. Fujimaki and S. Morinaga. Factorized asymptotic bayesian inference for mixture modeling. In *AISTATS*, volume 22, pages 400–408, 2012. 1, 2.2, 4, 4, 5.3
- [8] Z. Ghahramani and M. I. Jordan. Factorial hidden markov models. *Machine Learning*, 29:245–273, 1997. 10.1023/A:1007425814087. 1, 2.1, 3, 5.2
- [9] K. Hayashi and R. Fujimaki. Factorized asymptotic bayesian inference for latent feature models. In *27th Annual Conference on Neural Information Processing Systems (NIPS)*, 2013. 1, 2.2, 5.3, 5.3
- [10] D. Husmeier. Discriminating between rate heterogeneity and interspecific recombination in dna sequence alignments with phylogenetic factorial hidden markov models. *Bioinformatics*, 21(suppl 2):ii166–ii172, 2005. 1
- [11] H. Kim, M. Marwah, M. F. Arlitt, G. Lyon, and J. Han. Unsupervised disaggregation of low frequency power measurements. In *SDM*, pages 747–758, 2011. 1
- [12] J. Z. Kolter and T. Jaakkola. Approximate inference in additive factorial hmms with application to energy disaggregation. In *AISTATS*, volume 22, pages 1472–1482, 2012. 1
- [13] B. Leroux. Maximum-likelihood estimation for hidden markov models. *Stochastic processes and their applications*, 40(1):127–143, 1992. 5.4
- [14] C. A. McGrory and D. M. Titterington. Variational bayesian analysis for hidden markov models. *Australian & New Zealand Journal of Statistics*, 51(2):227–244, 2009. 2.1
- [15] G. Mysore and M. Sahani. Variational inference in non-negative factorial hidden markov models for efficient audio source separation. In *ICML*, 2012. 1
- [16] A. V. Nefian, L. Liang, X. Pi, X. Liu, and K. Murphy. Dynamic bayesian networks for audio-visual speech recognition. *EURASIP Journal on Advances in Signal Processing*, 2002(11):1274–1288, 2002. 1
- [17] L. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989. 2.1
- [18] S. T. Roweis et al. One microphone source separation. In *Advances in neural information processing systems*, pages 793–799. MIT, 1998, 2001. 1
- [19] J. Van Gael, Y. W. Teh, and Z. Ghahramani. The infinite factorial hidden markov model. In *Neural Information Processing Systems*, volume 21, 2008. 1, 2.1

A Appendix

Lemma 1. Suppose $\{z_{n,1}, \dots, z_{n,T_n}\}_{n=1}^N$ are N sequences of Bernoulli random variables, whose means are $\{p_{n,1}, \dots, p_{n,T_n}\}_{n=1}^N$. For the n -th sequence $\{z_{n,t}\}$, $z_{n,1}, \dots, z_{n,T_n}$ are independent with each other. Let $y_n = \sum_i z_{n,i}$, $\bar{y}_n = E[y_n] = \sum_i p_{n,i}$. Suppose further that $\bar{y}_n \rightarrow \infty$ as $T_n \rightarrow \infty$. Besides, there are N numbers $\{\hat{y}_n\}$, $\forall n, \hat{y}_n \approx \bar{y}_n$. When all T_n are large enough, the following bounds hold:

1)

$$\mathbf{E}[\log(y_n + 1)] = \log(\hat{y}_n + 1) + \frac{\bar{y}_n - \hat{y}_n}{\hat{y}_n + 1} + \epsilon_1;$$

2)

$$\mathbf{E}[y_n \log y_n] = \bar{y}_n \log \hat{y}_n + (\bar{y}_n - \hat{y}_n) + \epsilon_2.$$

Here ϵ_1, ϵ_2 are small bounded errors.

Proof.

1) Using the convexity of $-\log(y_n + 1)$, we easily obtain

$$\mathbf{E}[\log(y_n + 1)] \leq \log(\mathbf{E}[y_n] + 1) = \log(\bar{y}_n + 1). \quad (24)$$

We proceed to derive a lower bound of $\mathbf{E}[\log(y_n + 1)]$.

The following lower bound of the logarithm function is well known:

$$\log x > 1 - \frac{1}{x} \quad \text{for all } x > 0.$$

Substituting x with $\frac{y_n + 1}{\bar{y}_n + 1}$, we have

$$\log(y_n + 1) - \log(\bar{y}_n + 1) > 1 - \frac{\bar{y}_n + 1}{y_n + 1}.$$

After taking the expectation of both sides, it becomes

$$\mathbf{E}[\log(y_n + 1)] - \log(\bar{y}_n + 1) > 1 - (\bar{y}_n + 1)\mathbf{E}\left[\frac{1}{y_n + 1}\right]. \quad (25)$$

$\mathbf{E}\left[\frac{1}{y_n + 1}\right]$ is commonly referred to as the *negative moment* of y_n [3]. We apply the corollary in [3] that

$$\mathbf{E}\left[\frac{1}{y_n + 1}\right] = \int_0^1 G_n(t)dt, \quad (26)$$

where $G_n(t)$ is the probability generating function of y_n . It is known that $G_n(t) = \prod_{i=1}^{T_n} (q_{n,i} + p_{n,i} \cdot t)$, in which $q_{n,i} = 1 - p_{n,i}$. Thus (26) becomes

$$\mathbf{E}\left[\frac{1}{y_n + 1}\right] = \int_0^1 \prod_{i=1}^{T_n} (q_{n,i} + p_{n,i} \cdot t) dt. \quad (27)$$

We apply the *Inequality of arithmetic and geometric means* to $\prod_{i=1}^{T_n} (q_{n,i} + p_{n,i} \cdot t)$:

$$\begin{aligned} \forall t \geq 0 : \quad \prod_{i=1}^{T_n} (q_{n,i} + p_{n,i} \cdot t) &\leq \left(\frac{\sum_i q_{n,i}}{T_n} + \frac{\sum_i p_{n,i}}{T_n} t \right)^{T_n} \\ &= \left(1 - \frac{\bar{y}_n}{T_n} + \frac{\bar{y}_n}{T_n} t \right)^{T_n}. \end{aligned}$$

Let \bar{p}_n denote $\frac{\bar{y}_n}{T_n}$, and \bar{q}_n denote $1 - \frac{\bar{y}_n}{T_n}$. Then (27) becomes

$$\mathbf{E}\left[\frac{1}{y_n + 1}\right] \leq \int_0^1 (\bar{q}_n + \bar{p}_n t)^{T_n} dt = \frac{1 - \bar{q}_n^{T_n+1}}{(T_n + 1)\bar{p}_n} < \frac{1}{\bar{y}_n}. \quad (28)$$

Plugging (28) into (25), we have

$$\mathbf{E}[\log(y_n + 1)] - \log(\bar{y}_n + 1) > 1 - \frac{\bar{y}_n + 1}{\bar{y}_n} = -\frac{1}{\bar{y}_n}.$$

Thus $-\frac{1}{\bar{y}_n}$ is the lower bound of the approximation error, which tends to 0 as $T_n \rightarrow \infty$. Hence

$$\exists \epsilon_{1a} > 0, \text{s.t. } \forall T_n, \mathbf{E}[\log(y_n + 1)] > \log(\bar{y}_n + 1) - \epsilon_{1a}. \quad (29)$$

Combining (24) and (29), we have

$$|\mathbf{E}[\log(y_n + 1)] - \log(\bar{y}_n + 1)| < \epsilon_{1a}. \quad (30)$$

On the other hand, as $\hat{y}_n \approx \bar{y}_n$, the first order approximation of $\log(\bar{y}_n + 1)$ about \hat{y}_n is good:

$$\log(\bar{y}_n + 1) = \log(\hat{y}_n + 1) + \frac{\bar{y}_n - \hat{y}_n}{\hat{y}_n + 1} + \epsilon_{1b}, \quad (31)$$

where ϵ_{1b} is an error of order $o(\bar{y}_n^{-1})$ that tends to 0 when T_n increases.

Combining (30) and (31), we have

$$\left| \mathbf{E}[\log(y_n + 1)] - \log(\hat{y}_n + 1) - \frac{\bar{y}_n - \hat{y}_n}{\hat{y}_n + 1} \right| < |\epsilon_{1a}| + |\epsilon_{1b}|.$$

In other words, $\log(\hat{y}_n + 1) + \frac{\bar{y}_n - \hat{y}_n}{\hat{y}_n + 1}$ approximates $\mathbf{E}[\log(y_n + 1)]$ with a small bounded error. In addition,

this error tends to 0 quickly as the sequence length T_n increases.

2) It is easy to verify $y_n \log y_n$ is convex, and thus

$$\mathbf{E}[y_n \log y_n] \geq \bar{y}_n \log \bar{y}_n. \quad (32)$$

We proceed to derive an upper bound of $\mathbf{E}[y_n \log y_n]$.

Applying the inequality $\log x \leq x - 1$ by substituting x with $\frac{y_n}{\bar{y}_n}$, we have

$$y_n \log \frac{y_n}{\bar{y}_n} \leq y_n \left(\frac{y_n}{\bar{y}_n} - 1 \right) = \frac{y_n^2}{\bar{y}_n} - y_n. \quad (33)$$

Taking the expectation of both sides of (33), we have

$$\begin{aligned} & \mathbf{E}[y_n \log y_n] - \bar{y}_n \log \bar{y}_n \\ & \leq \frac{\mathbf{E}[y_n^2]}{\bar{y}_n} - \bar{y}_n \\ & = \frac{\text{Var}(y_n)}{\bar{y}_n} \\ & = \frac{\sum_i p_{n,i}^2}{\sum_i p_{n,i}} \leq 1. \end{aligned} \quad (34)$$

Combining (32) and (34), we have

$$|\mathbf{E}[y_n \log y_n] - \bar{y}_n \log \bar{y}_n| \leq 1. \quad (35)$$

Furthermore, when $\bar{y}_n \approx \hat{y}_n$, the first order approximation of $\log \bar{y}_n$ about \hat{y}_n is good, i.e.:

$$\begin{aligned} & \bar{y}_n (\log \bar{y}_n - \log \hat{y}_n) \\ & = \bar{y}_n (\bar{y}_n - \hat{y}_n) / \hat{y}_n + \epsilon_{2a} \\ & = (\bar{y}_n - \hat{y}_n) + \frac{(\bar{y}_n - \hat{y}_n)^2}{\hat{y}_n} + \epsilon_{2a}, \end{aligned} \quad (36)$$

where ϵ_{2a} is a small bounded error.

As $T_n \rightarrow \infty$, both \bar{y}_n and \hat{y}_n tend to ∞ . Therefore $\frac{(\bar{y}_n - \hat{y}_n)^2}{\hat{y}_n}$ is a small bounded error ϵ_{2b} . Then (36) becomes

$$\bar{y}_n \log \bar{y}_n - \bar{y}_n \log \hat{y}_n = (\bar{y}_n - \hat{y}_n) + \epsilon_{2a} + \epsilon_{2b}. \quad (37)$$

Combining (35) and (37), we have

$$\mathbf{E}[y_n \log y_n] = \bar{y}_n \log \hat{y}_n + (\bar{y}_n - \hat{y}_n) + \epsilon_2,$$

where $|\epsilon_2| \leq |\epsilon_{2a}| + |\epsilon_{2b}| + 1$ is also a small bounded error.

Corollary 2. $\{y_n\}$, $\{\bar{y}_n\}$ and $\{\hat{y}_n\}$ are defined as the same as in Lemma 1. Then the following bounds hold:

1)

$$\begin{aligned} & \mathbf{E}[\log \Gamma(y_n)] \\ & = \bar{y}_n \left(\log \hat{y}_n - \frac{1}{2\hat{y}_n} \right) - (\hat{y}_n + \frac{1}{2} \log \hat{y}_n) \\ & \quad + \frac{1}{2} (\log 2\pi + 1) + \epsilon_3; \end{aligned}$$

2)

$$\begin{aligned} & \mathbf{E} \left[\sum_n \log \Gamma(y_n) - \log \Gamma \left(\sum_n y_n \right) \right] \\ & = \sum_n \bar{y}_n \left[\log \left(\frac{\hat{y}_n}{\sum_m \hat{y}_m} \right) + \frac{1}{2 \sum_m \hat{y}_m} - \frac{1}{2 \hat{y}_n} \right] \\ & \quad + \frac{1}{2} \log \left(\sum_n \hat{y}_n \right) - \frac{1}{2} \sum_n \log \hat{y}_n \\ & \quad + \frac{1}{2} (N-1) (\log 2\pi + 1) + \epsilon_4. \end{aligned}$$

Here ϵ_3, ϵ_4 are small bounded errors.

Proof.

1) The Stirling's approximation for the log-Gamma function is:

$$\log \Gamma(y_n) = (y_n - \frac{1}{2}) \log y_n - y_n + \frac{1}{2} \log 2\pi + \epsilon_{3a}, \quad (38)$$

where ϵ_{3a} is an error term of $O(y^{-1})$, which goes to zero quickly as T_n increases. Therefore $\mathbf{E}[\epsilon_{3a}]$, denoted as $\bar{\epsilon}_{3a}$, is also a small bounded number.

Taking the expectation of both sides of (38), and applying Lemma 1, we obtain

$$\begin{aligned} & \mathbf{E}[\log \Gamma(y_n)] \\ & = \bar{y}_n \left(\log \hat{y}_n - \frac{1}{2\hat{y}_n} \right) - (\hat{y}_n + \frac{1}{2} \log \hat{y}_n) + \frac{1}{2} (\log 2\pi + 1) + \epsilon_3. \end{aligned}$$

2) Obtained by repeatedly applying 1), and combining terms involving \bar{y}_n .